

APS Security to Enable Remote Data Access

Rachana Ananthakrishnan, MCS

Brian Tieman, APS

1. Current usage scenario

A scientific group, comprising of a Principle Investigator (PI) and an identified set of other members in his group, apply for access to run experiments on the beam line for a set period of time. An evaluation is done and access is granted to selected scientific groups. Typically, access to the beam line is on order of days to a week. However, users generally need access to facility computational resources for as long as three to six months to completely process the data and move it to their home institutions.

Currently, data is acquired and stored in a local data repository under an account common to all the beam line users. All users performing experiments at a beam line have account level access to all data generated at that beam line while at the APS. In special cases, users may have a private account and auxiliary storage to which they may copy their data after acquisition. The security of the data relies solely on the propriety of individual users not to look at others data.

Once a group is provided access, it is granted physical access to the site and begins to run experiments during their allotted time period. Data acquisition software runs on the beam line with access to a local disk and writes data to the local disk at run time. The data is then staged out to an intermediate term (3-6 months) storage array on a machine at the APS.

The data remains on this storage array, under a common account, until it can be fully processed and moved to a user's home institution—typically via USB drives or anonymous FTP. While staged in this location, the data may be used in various ways at the user's discretion. During the user's allocated beam time, the data is often processed to provide quasi real-time feedback for optimizing the experiment in progress. Post beam time, the data is often reprocessed to optimize data quality and/or transferred to the user's home institution—typically via USB disks or an anonymous FTP server.

2. Security aspects of current usage scenario

- ◆ The process of granting access is outside the scope of this document.
- ◆ Enforcing granted access to the beam line is controlled by physical access to the site and paper work sent to operations.
- ◆ Data acquisition is not done as the user running the experiment, but rather as a shared account, per beam line. The data from each experiment is written to different directories/partitions on the local disk. The local disk directories and partitions are currently not protected by any file system permissions; hence the data acquired by one experiment is not protected from other users that run on the same beam line.
- ◆ There are currently no standard provisions for copying data from the local disk to specific user accounts. In special cases—typically frequent/heavy usage users—the data can be staged to user accounts. Often this is done with a simple copy command. The data that is

written to a user's account in intermediate term storage is protected by file system permissions.

- ◆ The intermediate term data storage machine is behind a firewall. Users requiring access to this storage from off-site are granted VPN accounts to access the data. Users not granted off-site permissions through the firewall rely on the resident staff to perform all required processing steps.
- ◆ Data is transferred off site either by a copy to portable USB disk drives or by an anonymous ftp server. A typical user's data volume for a one week run period is ~ 2TB.

3. Proposed usage scenarios

The proposed usage scenarios leverage following services that are assumed to be maintained by APS and provided at an endpoint accessible from outside the firewall:

1. APS Credential Service (ACS): This service provides an interface for an authenticated user to obtain short term PKI credentials and provisions remote client machines with PKI trust root information.
2. APS Data Transfer Service (ADTS): This service provides a remote endpoint to allow efficient and secure transfer of data to remote machines. It uses PKI security and authorization mechanisms to ensure secure access to data and can provide data integrity and privacy on transfer.
3. Technique based Scientific Portals (X-SP): These portals provide technique based (i.e. one portal per scientific technique independent of beam line: tomo-sp, 3dxdm-sp, xpc-sp, etc.) remote endpoint to allow efficient and secure utilization of local computational resources to perform experiments and process data in intermediate storage. It uses PKI security and authorization mechanisms to ensure secure access to computational resources.

3.1 Scenario 1: Secure local access

A scientific group SG1, with users Alice and Bob request access to the beam line for experiments and would like for the acquired data to be transferred to portable storage for them.

The group is granted access and a user account, say SG1, is created on the intermediate term storage. Not sure how best to do this...some sort of account creation/management service seems best. We may need to control which machines accounts get created on, etc... With the account information, they are also provided information on APS Credential Service (ACS) and client application downloads.

Comment [RA1]: Currently some users get ssh account on the intermediate term storage. can everyone get one?

Each user of SG1 logs onto the long term storage machine and uses the APS Credential Client installed to obtain a short term certificate from the ACS. Users will need to provide a user name/password to obtain short term certificates. This user name/password will be authenticated against LDAP or some other directory service. The user then runs GX-Map tool to set up authorization policy for the account they are logged in as. When this step is completed by Alice

Comment [RA2]: From Brian: I'm not sure how GX-Map works—we may need to hide this detail from average user.

and Bob, it sets policy on APS Data Transfer (ADTS) service to map both of them to the same account, SG1.

Now members of SG1 are ready to run experiments. On some day, Alice uses the beam line during the allotted time and runs the experiment. In addition to the acquisition software, she uses the APS Credential Client to contact the ACS to obtain a short term certificate. She uses the APS Data Transfer Client application on the beam line machine to start a transfer of the acquired data to intermediate term storage. The Data Transfer Client will require Alice's credential to authenticate with the ADTS, which then transfer's acquired data to account SG1. Once Alice has completed the experiment session, she deletes the data in the local disk and her credential in the local disk.

Similarly, Bob runs the experiment the subsequent day and he uses the APS Credential Client to obtain his short term credential. He then uses the APS Data Transfer Client to contact the ADTS to transfer data from his run. Since his credential is also mapped to account SG1 on the intermediate term storage, the data is transferred there. The ADTS logs the credential used to authenticate for the transfer and hence audit logs per user is maintained at the server. When he has completed the run, he deletes the data in the local disk and his credential.

Alice and Bob access the intermediate term storage locally and transfer the data to portable storage. Once they access window is completed, their accounts on the ACS and intermediate term storage is removed.

3.2 Scenario 2: Secure remote access

A scientific group SG2, with users Carol and Dave, request access to the beam line for experiments and would like for the acquired data to be transferred to remote machines in their home institutions.

The group is granted access and a user account, say SG2, is created on the intermediate term storage. With the account information, they are also provided information on APS Credential Service (ACS) and APS Data Transfer Service (ADTS).

Each user of SG2 then logs onto the intermediate term storage machine and uses the APS Credential Client installed to obtain a short term certificate from the ACS. Users will need to provide a user name/password to obtain short term certificates. This user name/password will be authenticated against LDAP or some other directory service. The user then runs GX-Map tool to set up authorization policy for the account they are logged in as. When this step is completed by Carol and Dave, it sets policy on APS Data Transfer (ADTS) service to map both of them to the same account, SG2.

Now the remote machine(s) to which data transferred is desired needs to be configured. A member of SG2 downloads and installs the APS Data Transfer Client and APS Credential Client on the identified remote machine(s). The clients are configured with the ADTS and ACS endpoints respectively.

With software configured, Carol now uses the beam line during the allotted time and runs the experiment. In addition to the acquisition software, she uses the APS Credential Client to contact the ACS to obtain a short term certificate. She uses the APS Data Transfer Client application on the beam line machine to start a transfer of the acquired data to intermediate term storage. The Data Transfer Client will require Carol's credential to authenticate with the ADTS, which then transfer's acquired data to account SG2. Once she has completed the experiment session, she deletes the data in the local disk and her credential in the local disk.

When the data in intermediate term storage is ready to be transferred to the remote machine, Carol uses the client software installed in the remote machine to initiate a transfer. She first uses the APS Credential Client to obtain a new short term credential and provision the remote machine with the required trust root information. She then uses the APS Data Transfer Client to download the data. The ADTS logs the user credential used for data download and hence audit logs per user access are stored at the server.

Once the experiment window is completed for group SG2, their accounts on ACS and intermediate term storage is disabled.

3.3 Scenario 3: Individual user in multiple groups

Alice and Carol from previous scenarios are part of yet another scientific group and request access to the beam line for experiments. They would like for the acquired data to be transferred to remote machines in their home institutions.

The group is granted access and a user account, say SG3, is created on the intermediate term storage. With the account information, they are also provided information on APS Credential Service (ACS) and APS Data Transfer Service (ADTS).

Alice and Carol can leverage the software set up on their remote machines without any additional configuration. They don't need not require credentials from the ACS, but rather reuse the existing credential. But they will need to log in as SG3 and run the GX-Map tool to set up authorization policy to access data in SG3.

When Alice runs experiment on the beam line machine for this group SG3, she runs the ACS client to retrieve new credentials. When using the APS Data Transfer Clients to initiate transfer to storage machine, she chooses her username as "SG3". This ensures that the data being transferred is stored in the SG3 account. She then continues with the process like in the previous scenario. When the experiment is completed, she deletes data and credential from local disk

To transfer the acquired data to a remote machine, Carol uses the same clients installed and the process. When she initiates the data transfer from remote machine, she again chooses the local account corresponding to the scientist group, SG3 in this case.

3.4 Scenario 4: Secure remote access to Scientific Portal

A scientific group SG4, with users Eva and Fred, request remote access to the beam line for 3DxDM experiments and post-experiment data processing and would like for the acquired data to be transferred to remote machines in their home institutions.

The group is granted access and a user account, say SG4, is created on the intermediate term storage. With the account information, they are also provided information on APS Credential Service (ACS) and APS Data Transfer Service (ADTS) and the 3D X-Ray Micro Diffraction Scientific Portal (3DXDM-SP).

Each user of SG4 then logs onto the intermediate term storage machine and uses the APS Credential Client installed to obtain a short term certificate from the ACS. Users will need to provide a user name/password to obtain short term certificates. This user name/password will be authenticated against LDAP or some other directory service. The user then runs GX-Map tool to set up authorization policy for the account they are logged in as. When this step is completed by Eva and Fred, it sets policy on APS Data Transfer (ADTS) service to map both of them to the same account, SG2.

Now the remote machine(s) to which data transferred is desired needs to be configured. A member of SG4 downloads and installs the 3DXDM-SP client. The client is configured with the ADTS and ACS and 3DXDM-SP endpoints respectively.

With software configured, Eva now uses the beam line/computational resources during the allotted time and runs the experiment. She uses the APS Credential Client to contact the ACS to obtain a short term certificate. She uses the 3DXDM-SP to control the experiment and computational resources. The 3DXDM-SP will require Eva's credential to allow control of the resources. She uses the APS Data Transfer Client application on the beam line machine to start a transfer of the acquired data to intermediate term storage. The Data Transfer Client will require Eva's credential to authenticate with the ADTS, which then transfer's acquired data to account SG4. Once she has completed the experiment session, she deletes the data in the local disk and her credential in the local disk.

When the data in intermediate term storage is ready to be transferred to the remote machine, Eva uses the client software installed in the remote machine to initiate a transfer. She first uses the APS Credential Client to obtain a new short term credential and provision the remote machine with the required trust root information. She then uses the APS Data Transfer Client to download the data. The ADTS logs the user credential used for data download and hence audit logs per user access are stored at the server.

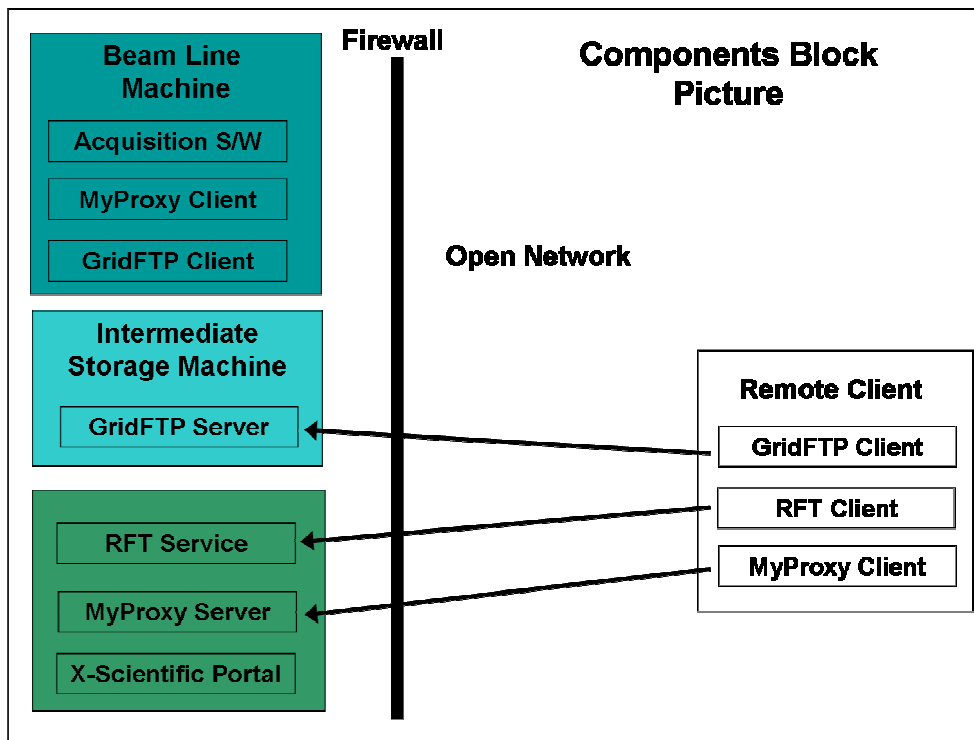
Once the experiment window is completed for group SG4, their accounts on ACS and intermediate term storage is disabled.

4. Requirements from proposed usage scenarios

- ◆ Allow for transfer of the acquired data to a remote location of user's choice.

- ◆ Protect experiment data such that access is restricted to only scientific group participating in that experiment
- ◆ Low overhead mechanism to issue credentials to users
- ◆ Light weight user clients for credentials management and data download
- ◆ Automated trust root provisioning of client machines
- ◆ Data encryption as an option
- ◆ Access policy on scientific group, rather than per user.
- ◆ Track individual user access for audit purposes
- ◆ Support for individual users belonging to multiple groups operating at the same beam line.

5. Proposed Deployment



Services run by APS and hosted at an endpoint that is accessible from outside the firewall:

1. APS Credential Service (ACS): This is an installation of MyProxy Online CA service and issues short term credentials to authenticated users. The authentication system could leverage the current user accounts scientific groups are provided when they are granted access to the beam line. Alternatively registration systems like PURSe, which is portal-based user registration service, maybe installed to create and manage user accounts.

With either authentication mechanism, MyProxy can use a PAM module to leverage the back end system and issue new credentials.

2. APS Data Transfer Service (ADTS): This is a GridFTP Server that uses the GridFTP protocol to provide secure, robust, fast and efficient transfer of (especially bulk) data. The service is set up to require PKI authentication and is configured with authorization policy to restrict access to the data. Alternatively, this could be an endpoint to Globus Reliable File Transfer Service (RFT) which provides an additional reliability layer by moving large data sets on behalf of the client, without the client requiring maintaining active sockets with the GridFTP server. RFT also provides the fail over and recoverability, in some cases, by maintaining state of the transfer in a database.
3. Scientific portal Services (X-SP): These are APS developed services that users interact with to workflow experiments. The services are set up to require PKI authentication. A secondary authorization service may be developed and installed to avoid conflicts for “one of” resources such as beam lines.

Software installed on beam line machines:

1. Data Transfer Client Application: This would be a combination of GridFTP client software and some scripts that can feed information about new data written by acquisition software.
2. APS Credential Client Package: This would be MyProxy Logon Clients that allows a client to obtain short term credential from the MyProxy server and also provision the local machine with trust root information. The package will also include the MyProxy Destroy command that allows a user to remove the user’s credential in the local disk.
3. Scientific Portal Client (X-SP): This would be a customized client for an experimental technique. This client may integrate the Data Transfer Client, APS Credential Client and specific experiment control clients into one integrated package for control of experiments throughout the experiment life-cycle.

Software for intermediate term storage access:

1. APS Data Transfer Service: A GridFTP server installation, configured to trust the ACS CA and GridMap file authorization. The transfer server should not allow anonymous or username/password access to the data.
2. gx-map: This software need to be installed to allow clients who log on to the system to add themselves to the GridMap file. Explore gx-map option without shared file system, could not find too many details.

Software for client downloads:

1. Data Transfer Client application

2. Credential Service Client application

6. Tasks:

1. Prototype 1: This will enable remote data transfer only and assume each user is only part of a single group. Installations will include GridFTP server a machine outside the firewall with GridMap file, gx-map on the storage cluster machine and a MyProxy server outside the firewall. MyProxy server is hooked into a simple back end database system with username/password (if there is an existing system we can use, we should.)
2. Prototype 2: This will allow for users to be part of multiple groups. Data Transfer client package enhanced with ability to select the local user account to transfer data from is provided as part of client downloads. This requires the user to enter the storage machine username associated with the scientific group they need to access data for.
3. Prototype 3: This will secure the data on the local disk at the beam line. Installation will include GridFTP client libraries on the beam line machine that runs Windows and scripts that can leverage the polling capability to transfer acquired data immediately to storage machine. Also, MyProxy client library to obtain user credentials and a script that allows user to delete all local data.
4. Prototype 4: This will allow for reliable transfer of remote data. Installation will include an RFT server on a machine outside the firewall and client downloads of RFT client libraries. RFT should be configured with Secure Transport access and GridMap authorization. Can RFT pass chosen local username?

6. References

- ◆ Globus GridFTP Server <http://www.globus.org/toolkit/docs/development/4.2-drafts/data/gridftp/index.html>
- ◆ Globus RFT <http://www.globus.org/toolkit/docs/development/4.2-drafts/data/rft/index.html>
- ◆ MyProxy Server <http://www.globus.org/toolkit/docs/development/4.2-drafts/security/myproxy/index.html>
- ◆ Gx-map <http://users.sdsc.edu/~gxmap/>
- ◆ PURSe <http://dev.globus.org/wiki/Incubator/PURSe>

7. Things to add:

- ◆ Installation links
- ◆ Firewall requirements/issues
- ◆ PURSE for registration
- ◆ Time frame
- ◆ Separate document? Computer cluster access: potentially transfer data to storage and submit jobs to WS GRAM on some compute cluster and feed results back to the user. I

have a working prototype of a Scientific Portal client. The client can create simple workflows of services such as RFT and in house developed Introduce and gRavi services for control of computational resources including our cluster. I do not yet have GRAM installed, but I do have a gRavi service that can fire up our cluster. I actually have two sets of services—one set allows for interactive data browsing of samples within the system and one set allows for pipelining samples through a workflow. I'd like to see this integrated into the security plan—hence the addition of scenario 3.

Questions

1. Is the beam line machine on any shared file system? Today if copy is done from local disk to storage array, what is the performance on this? This is an attempt to establish motivation for GridFTP.
2. Currently some users get ssh accounts on machine with storage array. Can everyone be given an account that?

Document draft date: 06/11/2008